How hard is it to control an election by breaking ties?

Nicholas Mattei, Nina Narodytska, and Toby Walsh
NICTA and UNSW
Sydney, Australia
{nicholas.mattei, nina.nardytska, toby.walsh}@nicta.com.au

Abstract

We study the computational complexity of the problem of controlling the result of an election by breaking ties. When the chair is only asked to break ties to choose between one of the co-winners, the problem is trivially easy. However, in multi-round elections like STV, we prove that it can be NP-hard for the chair to compute how to break ties to ensure a given result. Our results contain several surprises. For example, whilst it is NP-hard to compute a manipulating vote for a multi-round rule like Nanson, it is polynomial for the chair to control the result by breaking ties. As a second example, it can be NP-hard to control an election by breaking ties even with a simple two-stage voting rule.

1 Introduction

Voting is a general mechanism to combine together orderings (e.g. preferences of agents over different plans, or rankings of web pages by different search engines). One concern is that the chair may manipulate the result. For example, the chair might introduce a spoiler candidate or delete some votes. Bartholdi, Tovey and Trick [1992] explored an interesting barrier to such manipulation; perhaps it is computationally too difficult for the chair to work out how to perform such control? They proved that many types of control problems are NP-hard even for simple voting rules like plurality.

Interestingly, one type of control not considered by Bartholdi, Tovey and Trick is control by choosing how ties are broken. This is surprising since the chair is actually the person who breaks ties in many elections. For example, the Speaker in many parliaments has the casting vote in case of a tied vote. Another reason to consider such control is that in many other elections the tie breaking rule is unspecified or is left ambiguous. The chair therefore has an opportunity to influence the outcome. This control problem also avoids one of the criticisms raised against the analysis of some of the other forms of control. In particular, many complexity results about control suppose that the chair has complete knowledge of the votes. This might be considered unreasonable. For example, how do we know how voters will rank a new spoiler candidate till their candidature has been announced? When studying control by breaking ties, it is natural to suppose the chair knows how the votes are cast when asked to break a tie.

Control by tie-breaking is equivalent to the problem of determining if a chosen alternative can win under *some* tie-breaking rule. This is called parallel universe tie-breaking (PUT) [Conitzer *et al.*, 2009] and has been used to create neutral rules [Tideman, 1987]. Deciding if a candidate is the winner of such a neutral rule with ranked pairs voting has recently been shown to be NP-complete [Brill and Fischer, 2012]. It immediately follows that control by tie-breaking for ranked pairs is NP-complete.

Tie-breaking has played an important role in some of the earliest literature on computational social choice. For example, Bartholdi, Tovey and Trick [1989] proved that a single agent can manipulate a Copeland election in polynomial time when ties are broken in favour of the manipulators, but manipulation becomes NP-hard with the tie-breaking rule used in chess competitions. With Copeland voting, Faliszewski, Hemaspaandra and Schnoor [2008] proved that the choice of how ties are scored can change the computational complexity of computing a manipulation from polynomial to NP-hard. More recently work by Obraztsova, Elkind and Hazon [2011] and Aziz et al. [2012] considered the impact of different randomized tie-breaking schemes on the computational complexity of computing a manipulation. They proved, for example, that all scoring rules are polynomial to manipulate for some tie-breaking rules but not others; additionally rules like maximin, STV and ranked pairs are NP-hard.

In this paper, we study the computational complexity of control by breaking ties. We show that when the chair only breaks ties to choose between co-winners, as is the case in many single round rules, control by tie-breaking is polynomial. On the other hand, for many multi-round rules like STV and Coombs, the chair may have to break multiple ties, and the control by tie-breaking problem is NP-complete. Even with two-stage rules, where the chair may have to break ties only twice, the control by tie-breaking problem can be NP-complete. We also consider tournament rules like Copeland where tie-breaking is already known to have a computational impact on strategic voting.

2 Formal background

A *profile* is a set of n total orders (votes) over m candidates. A *voting correspondence* is a function mapping a profile onto

a set of *co-winners*. A tie-breaking rule can then be used to return the unique winner. Let $N(i, j)$ be the number of voters preferring i to j. We consider some of the most common voting correspondences.

Scoring rules: $(w_1, \ldots, w_m)$ is a vector of weights where the ith candidate in a vote scores w_i, and the co-winners are the candidates with highest total score. *Plurality* has $w_1 = 1$, and $w_i = 0$ for $i > 1$; *veto* has $w_i = 1$ for $i < m$, and $w_m = 0$, ; k-*approval* has $w_i = 1$ for $i \leq k$, and $w_i = 0$ for $i > k$; *Borda* has $w_i = m - i$.

Plurality with runoff: If one candidate has a majority, then she wins. Otherwise we eliminate all but the two candidates with the most votes and apply the plurality rule.

Black's: If one candidate is the Condorcet winner, a candidate preferred by a majority of voters to all others, then she wins. Otherwise, we apply the Borda rule.

Bucklin: This rule proceeds in rounds. The winning round is the smallest value k such that the k-approval score of at least one candidate exceeds $\lfloor n/2 \rfloor$. The Bucklin score of a candidate is his k-approval score, where k is the winning round. The co-winners are the candidates with the largest Bucklin score. The *simplified Bucklin* procedure is the same except that all candidates with score exceeding $\lfloor n/2 \rfloor$ are co-winners.

Fallback: This is a combination of Bucklin and approval voting. Voters approve a subset of candidates and rank this subset. If there is a k such that the k-approval score of at least one candidate, the sum of the approvals appearing in the first k places of each voter's ranked order, exceeds $\lfloor n/2 \rfloor$ then the co-winners are the set of candidates exceeding this threshold. If there is no such k (as no candidate receives enough approvals), the winner is the approval winner.

Single Transferable Vote (STV): This rule requires up to $m - 1$ rounds. In each round, the candidate with the least number of voters ranking him first is eliminated until one of the remaining candidates has a majority.

Nanson and Baldwin: These are elimination versions of Borda voting. In each round of Nanson, we eliminate all candidates with less than the average Borda score. In each round of Baldwin, we eliminate the candidate with the lowest Borda score.

Coombs: This is the elimination version of veto voting. In each round, we eliminate the candidate with the lowest veto score until we have one candidate with a plurality score of $n/2$ or greater. In the simplified version of Coombs, we eliminate the candidate with the lowest veto score until one candidate remains.

Cup: Given a schedule T and a labelling L, we run a knockout tournament. Candidates are compared pairwise with the winner in each competition going forwards to the next round.

Copeland$^\alpha$: The candidates with the highest Copeland$^\alpha$ score win. The Copeland$^\alpha$ score of candidate i is $\sum_{j \neq i} (N(i, j) > \frac{n}{2}) + \alpha \cdot (N(i, j) = \frac{n}{2})$. In the second order Copeland rule, if there is a tie, the winner is the candidate whose defeated competitors have the largest sum of Copeland scores.

Ranked pairs: We consider all pairs of candidates in order of the pairwise margin of victory, from greatest to least. For each considered pair, we construct an ordering which ranks these candidates unless it creates a cycle. The winner is the candidate at the top of the constructed ordering. For a neutral variant, the co-winners are any candidate who can be made top element under some tie breaking order [Brill and Fischer, 2012].

Maximin: The Maximin score of a candidate is the number of votes received in his worst pairwise election. The co-winners are the candidates with the largest such score.

We consider the following decision problems. In each, we are given a profile consisting n votes which are strict linear orders over a set of m candidates, a preferred candidate p and a voting correspondence. We use $a > b$ to indicate that candidate a is strictly preferred to candidate b and a b to indicate that two candidates are tied. In the *control by tie-breaking problem*, we wish to decide if the chair can break ties to make p the unique winner. Note that the control by tie-breaking problem is equivalent to the *possible winners problem* when the tie-breaking rule is unknown.

A voting rule is *vulnerable* to such control if this problem is polynomial, and *resistant* if it is NP-hard. Additionally, a rule can be *immune* to a form of control if there exists no profile where the given form of control is effective, conversely a rule is *susceptible* there exists a profile that can be changed.

In the *manipulation problem*, we wish to decide if we can cast one additional vote to make p win. All our results here apply to the variants of the manipulation problem in which we break ties in favour of or against the manipulator. Finally, in the *manipulation problem with random tie-breaking*, we are also given a probability t and we wish to decide if we can cast one additional vote to make p the winner with probability at least t supposing ties are broken uniformly at random between candidates.

3 Relationship to manipulation

We start by considering how control by breaking ties is related to other manipulation problems. A little surprisingly, the complexity of control by breaking ties is not related to that of the manipulation problem with random tie-breaking.

Theorem 1. *There exists a voting correspondence such that the control by tie-breaking problem is polynomial but the manipulation problem with random tie-breaking is NP-complete (and vice versa).*

Proof: In Theorem 13, we prove that the control by tie-breaking problem for Copeland is polynomial. On the other hand, [Obraztsova and Elkind, 2011] prove that the manipulation problem with random tie-breaking for Copeland is NP-complete.

Consider the voting rule that eliminates half the candidates using the veto rule, then elects the plurality winner. In Theorem 15, we prove that the control by tie-breaking problem for this rule is NP-complete. However, the manipulation problem with random tie-breaking for this rule is polynomial since we

can exhaustively try all $m(m-1)/2$ votes with different candidates in the first and the last position, and compute the probability that our distinguished candidate wins in each case. □

The computational complexity of control by breaking ties is also not related to that of the manipulation problem (where ties are broken in a fixed order, in favour or against the manipulator).

Theorem 2. *There exists a voting correspondence such that the control by tie-breaking problem is polynomial but the manipulation problem is NP-complete (and vice versa).*

Proof: In Theorem 9, we prove that the control by tie-breaking problem for Nanson is polynomial. On the other hand, the manipulation problem for Nanson is NP-complete [Narodytska *et al.*, 2011].

Consider again the voting rule that eliminates half the candidates using the veto rule, then elects the plurality winner. In Theorem 15, we prove that the control by tie-breaking problem for this rule is NP-complete. However, the manipulation problem is polynomial since we can exhaustively try all $m(m-1)/2$ votes with different candidates in first and last position. □

These results indicate that we must search for other properties that may link the single manipulator problem with the control by tie-breaking problem.

4 Breaking ties once or twice

We start with some very simple cases. When tie-breaking only ever takes place once and at the end, then the chair is choosing between the co-winners. In such cases, control by breaking ties is trivially polynomial. The chair can ensure a candidate p wins if and only if p is amongst the co-winners.

Theorem 3. *The control by tie-breaking problems for scoring rules, Bucklin and maximin are polynomial.*

We move to some slightly more complex cases. In plurality with runoff, the chair breaks ties at most twice: once to break ties between candidates for the runoff, and once to break ties in the runoff. Control by breaking ties remains polynomial.

Theorem 4. *The control by tie-breaking problem for plurality with runoff is polynomial.*

Proof: $O(m)$ possible candidates can enter the runoff with the distinguished candidate. We can try all possibilities. □

Another rule that is vulnerable to control by breaking ties is Black's rule.

Theorem 5. *The control by tie-breaking problem for Black's rule is polynomial.*

Proof: There cannot be a tie in the first stage of Black's rule. Either we have a Condorcet winner, who is then elected, or we do not. If we move to the next stage we only need to break ties at the end of the rule between candidates with equal top Borda score. □

Another rule that is vulnerable to control by breaking ties is fallback voting. This is perhaps a little surprising as it holds the current record of resistance to 20 out of the 22 standard methods of control [Rothe and Schend, 2012].

Theorem 6. *The control by tie-breaking problem for fallback voting is polynomial.*

Proof: The chair only ever needs to tie-break once between a set of co-winners. □

To have any resistance to control by tie-breaking, we need more complex tie-breaking. One place to see more complexity is with multi-round rules like STV and Coombs in which candidates are successively eliminated. Such rules increase the number of times ties may need to be broken.

5 Multi-Round Voting Rules

We now move to multi-round voting rules. Bartholdi and Orlin [1991] showed that the manipulation problem for STV is NP-complete; Conitzer et al. [2009] showed that the winner determination under PUT for STV, and therefore control by tie-breaking, is NP-complete.

Theorem 7 (Conitzer et al. [2009])**.** *The control by tie-breaking problem for* STV *is NP-complete.*

5.1 Nanson's and Baldwin's Rules

We next consider Nanson and Baldwin's voting rules. These are multi-round rules that successively eliminate candidates based on their Borda score. The manipulation problem for Baldwin's rule is NP-complete. Baldwin's rule is also resistant to control by breaking ties.

Theorem 8. *The control by tie-breaking problem for Baldwin's rule is NP-Complete.*

Proof: We will modify the initial scores of the candidates in the NP-completeness proof of Baldwin manipulation [Narodytska *et al.*, 2011]. The chair will set the tie-breaking order such that we select exactly a subset of sets that give us an exact cover in an instance of EXACT COVER BY 3 SETS (X3C). Given two sets $V = \{v_1, \dots, v_q\}$, $q = 3t$, and $S = \{S_1, \dots, S_t\}$, where $t \geq 2$ and for all $j \leq t$, $|S_j| = 3$, and $S_j \subseteq V$ we create an instance with the set of candidates $C = \{p, d, b\} \uplus V \uplus A$. Note that p is the preferred candidate, members of $A = \{a_1, \dots, a_t\}$ correspond to the 3-sets in S, and $m = |C| = q + t + 3$. The construction is made up of two parts. The first set of votes P_1 remains unchanged from [Narodytska *et al.*, 2011]. The second set of votes P_2, which are the votes that set the initial score differences between the candidates, are modified so that the dangerous candidates are tied with p (rather than having one more vote than p).

With the modified scores, all candidates in A are tied with p in the first round. Since all these candidates are tied, the tie-breaking rule must choose one to remove in each round. In round $4k = 0, \dots, q/3$ we must select some set of candidates $a_1, \dots, a_{q/3}$ to eliminate (in the interleaving rounds $4k+1, 4k+2, 4k+3$, the elements v_i in the set S_j corresponding to a_j will drop out). At each $4k$, p will be tied with the remaining candidates in A which correspond to sets S until there are no more sets to cover (after $4q/3$ rounds). Then p will be tied with b if and only if we have eliminated a cover and the remaining a_j that were not part of the cover. We then select p to win over b and the rest of the proof proceeds as in [Narodytska *et al.*, 2011]. □

Nanson's rule is another multi-round rule based on successively eliminating candidates according to their Borda scores. The manipulation problem for Nanson's rule is NP-complete

[Narodytska *et al.*, 2011]. Surprisingly, Nanson's rule is vulnerable to control by breaking ties.

Theorem 9. *The control by tie-breaking problem for Nanson's Rule is polynomial.*

Proof: In Nanson's rule, we eliminate all alternatives with less than the average Borda score. The only time that we break ties is in the final round when multiple candidates have the maximal Borda score. □

5.2 Coombs Rule

We finish our consideration of multi-round rules with Coombs rule which successively eliminates the candidate with the largest number of last place votes.

Theorem 10. *The control by tie-breaking problem for Coombs rule is NP-complete.*

Proof: The result holds for both the simplified and unsimplified Coombs rule. We start from a reduction used to show the manipulation problem for Coombs rule is NP-complete. Given an instance of X3C and the the profile E defined in Theorem 3 from Davies et al. [2012] we prove hardness of the control by tie-breaking problem with a slight modification of the scores: we increase the initial veto scores of s_2 and d_0 by 1 each. The profile E, generally, creates a voting instance where a cover is selected and then verified through sequential eliminations through a complex setting of initial candidate scores. Here, we need to show that the influence that a single manipulator has on the outcome of the election can be again simulated by the tie-breaking rule.

During the first $4m$ rounds, the manipulator can *only* change the outcome of E at rounds $p \in \{1, 5, 9, \ldots, 4m\}$, where exactly two candidates are tied. Therefore, we can simulate the manipulator's influence using the tie-breaking rule. The second point where a manipulator's vote has an influence is at the third stage after $d_1, \ldots, d_n$ are eliminated (in this case there is a cover). Namely, the manipulator uses 1 veto-point to tie c and d_0. This allows d_0 to be eliminated. However, as we increased the initial veto-score of d_0, d_0 and c are tied after $d_1, \ldots, d_n$ are eliminated. Hence, by ranking c above d_0, the manipulator can eliminate d_0. Similarly, after $d_0, d_1 \ldots, d_n$ are eliminated, the score of c is the same as the score of s_2 as we increased the initial veto-score of s_2 by 1. Hence, by means of the tie-breaking rule, s_2 can be eliminated before c. The elimination order at the 4th stage is independent of the manipulator's vote. Hence, p wins the election if and only if we select a cover during the first $4m$ rounds by means of breaking ties appropriately. Hence, determining if a preferred candidate is a winner of the election under by breaking ties appropriately is NP-complete. □

6 Tournament-style rules

Like multi-round rules, tournament[1] based rules like Cup, Copeland, and ranked pairs may require multiple ties to be resolved.

[1] We use tournament here in it's common language meaning as a sports or matchup competition; not necessarily in its mathematical sense as a complete directed graph.

6.1 Cup

In order to determine the tie breaking order, we simply run the algorithm from Theorem 2 in [Conitzer *et al.*, 2007] that computes a manipulating vote and use the returned manipulation as the linear order for tie-breaking.

Theorem 11. *The control by tie-breaking problem for the Cup rule is polynomial when each candidate appears only once in T.*

Notice that the above procedure always returns a linear tie-breaking order. When each candidate appears only once, a manipulator cannot benefit by breaking ties with an order that violates transitivity. To profit from a non-transitive order, we would need multiple pairwise comparisons between candidates. For example, in double elimination tournaments, candidates appear twice. In fact, we only need a linear schedule and one candidate to appear twice to see the difference.

Example. *Consider the majority relation shown below which we can realise through votes according to McGarvey's Theorem [1953]. All un-drawn arrows go from left to right. The relation between a, b, and c is unspecified (tied).*

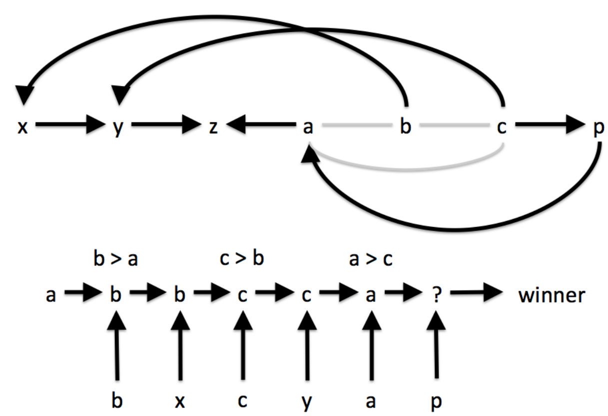

For the given agenda we must select the pairwise relations between a, b, and c. Notice that the only way to ensure that p wins the tournament is to have $c > b$, $b > a$ and $a > c$.

If we allow candidates to enter the tournament more than once, and if the tournament can have arbitrary shape, the control by tie-breaking problem becomes hard for both transitive and non-transitive tie-breaking orders.

Theorem 12. *When the Cup schedule T can have arbitrary shape and candidates can appear more than once, control by tie-breaking is NP-complete.*

Proof: (Sketch) Reduction from 3SAT. Given a set of clauses $K = \{k_1, \ldots, k_n\}$ and a set of literals with their negations $L = \{l_1, \overline{l_1}, \ldots, l_m, \overline{l_m}\}$ we create a Cup tournament with candidates $C = K \uplus L \uplus \{p\}$. The majority relation over the candidates has p defeating all elements of L; each literal and its negation are tied and defeating any literal with a higher number ($l_1 \sim \overline{l_1} > l_2 \sim \overline{l_2}$); each clause is defeated by only those literals that would satisfy it while defeating all other literals and p. We construct T as follows: for each k_i we pair the three literals that would satisfy k_i with their negations in a match-up. Then, k_i plays the winner

of these three match-ups, sequentially. The output of these smaller tournaments all play p sequentially. If there is a tie-breaking order which selects either l_1 or $\overline{l_1}$ for every literal such that all clauses are satisfied, then p will win the Cup. Each k_i will face each literal or negation in k_i, depending on the tie-breaking rule. Since only literals that satisfy k_i defeat it, one must be selected, otherwise p will lose to k_i when p plays the winner of each of the sub-tournaments. If k_i is satisfied by one (or more) of its literals, then k_i will be eliminated and a literal will face (and lose to) p in the latter part of the tournament. Hence p will win the tournament if and only if there is tie-breaking rule that satisfies the 3SAT instance. $\Box$

6.2 Copeland

For the Copeland rule, we know that the second-order tie-breaking rule is NP-hard to manipulate [Bartholdi *et al.*, 1989]. We can also devise other tie-breaking rules to add to Copeland to make manipulation NP-hard [Obraztsova *et al.*, 2011]. On the other hand, the regular Copeland rule is vulnerable to control by breaking ties. We merely need to break ties in favour of the preferred candidate, and for other pairs of candidates, we use the algorithm of Bartholdi, Tovey and Trick [1989] to minimize the Copeland scores of all other candidates.

Theorem 13. *The control by tie-breaking problem for Copeland is polynomial.*

If we permit the tie-breaking rule to be non-transitive, then we can increase the possibilities for manipulation. Specifically, allowing the tie-breaking rule to be non-transitive increases the potential for control of the tie-breaking rule under Copeland. Consider the election with 6 votes: (d, f, g, p, a, b, c), (d, f, g, p, c, b, a), $2 \times (p, a, b, c, d, f, g)$, and $2 \times (c, b, a, d, f, g, p)$. Suppose we want p to win. The Copeland scores are $p = a = b = c = d = 3, f = 2, g = 1$. We need to submit a tie-breaking order that will resolve the pairwise ties between a, b and c. There is no transitive order that we can submit to resolve these such that p wins. However, we can submit pairwise preferences that will maintain the cycle and allow p to win.

We note that such tie-breaking is close to the problem of manipulating a Copeland election with *irrational* voters [Faliszewski *et al.*, 2009] which is polynomial time computable [Faliszewski *et al.*, 2008]. If we allow the chair to specify the result of each pair-wise tie separately (and thus to break ties non-transitively), Copeland remains vulnerable to control by breaking ties.

Theorem 14. *The control by tie-breaking problem for the Copeland rule is polynomial, even when tie-breaking is specified in terms of a perhaps non-transitive ordering on pairwise contests.*

The Copeland rule offers another interesting control opportunity for the chair. The chair might be in a position to set α, the score that a candidate receives in the event of a tie in the tournament graph. The choice of α has an impact on the computational complexity of computing a manipulation [Faliszewski *et al.*, 2008; Faliszewski *et al.*, 2009]. However, the control problem of choosing the best value of α to make a given candidate win is polynomial.

7 Two stage rules

We have seen that when we break ties only once to decide between the co-winners, voting rules are vulnerable to control by breaking ties. We have also seen that when we break ties multiple times, some multi-round voting rules like STV are resistant to control by breaking ties. This leaves open the question of whether we ever have resistance to control by breaking ties when the chair only has a small, but fixed number of opportunities to break ties. We consider here two stage rules where the chair only has at most two opportunities to break ties. Surprisingly, control by breaking ties can be intractable in this case.

Theorem 15. *There exists a two stage voting rule that combines veto and plurality voting for which the control by tie-breaking problem is NP-complete.*

Proof: We consider the rule that first eliminates half the candidates using the veto rule, then elects the plurality winner. Clearly, the control problem is in NP. We merely need to give the subset of candidates selected after tie-breaking to go through to the runoff. To show NP-hardness, we adapt the reduction from X3C used in the proof of Theorem 3 in [Bartholdi *et al.*, 1992] that demonstrates control by elimination of candidates for plurality is NP-hard. This reduction uses $n + 4m/3 + 2$ candidates where m is the size of the set being covered, n is the number of 3-element sets from which the cover is built. We double the number of candidates to $2n + 8m/3 + 4$ with $n + 4m/3 + 2$ additional dummy candidates d_i that occur in the same fixed order in every vote. The first $n + m + 1$ candidates appear at the front of the votes, whilst the last $m/3 + 1$ appear at the end. Now the X3C problem has a solution if and only if we can remove $m/3$ candidates from the original votes to ensure the distinguished candidate is the plurality winner. With our two stage rule, one of the dummy candidates has all the vetoes so will be eliminated. The chair therefore has to tie-break on which of the other $2n + 8m/3 + 3$ candidates must be eliminated. To ensure that the distinguished candidate is the plurality winner, the chair's tie-breaking must eliminate all $n + m + 1$ dummy candidates at the front of the vote, plus $m/3$ of the candidates from the original election corresponding to the cover. Hence, the X3C problem has a solution if and only if the chair can tie-break to ensure the distinguished candidate wins. $\Box$

8 Hybrid voting rules

Conitzer and Sandholm [2003] give a general construction that builds a two-stage voting rule that often makes it intractable to compute a manipulating vote. This construction runs one round of the Cup rule, eliminating half of the candidates, and then applies the original base rule to the candidates that remain. For the base rule X, we denote this as $Cup_1 + X$. The control by tie-breaking problem is also typically intractable for such two-stage voting rules.

Theorem 16. *The control by tie-breaking problem for $Cup_1 + Plurality$, $Cup_1 + Borda$, and $Cup_1 + Maximin$ are NP-complete.*

Proof: Consider the reduction from SAT used in Theorem 2 in Conitzer and Sandholm [2003] showing that it is NP-hard

to construct a single vote to ensure a distinguished candidate wins $Cup_1 + Plurality$. This reduction sets up a profile in which the candidates c_{+v} and c_{-v} corresponding to a literal and its negation which are paired in the first round of Cup are tied. There is a vote that breaks these ties so that the distinguished candidate wins if and only if the SAT instance is satisfiable. Let us consider just the original profile, without the single manipulating vote. Now, the chair can break these ties so that the distinguished candidate wins if and only the SAT instance is satisfiable. The other proofs are similar and are adapted from the corresponding reductions in [Conitzer and Sandholm, 2003]. □

Elkind and Lipmaa [2005] generalize this construction to run some number of rounds, k, of one rule (not necessarily Cup) before calling a second rule. This again makes computing a manipulating vote NP-hard in many cases. Control by tie-breaking for such hybrids is often NP-hard as tie-breaking can simulate the manipulating vote used in the proofs in Elkind and Lipmaa [2005]. For example, control by tie-breaking for $\textsc{Hyb}(\text{STV}_k, Y)$ and $\textsc{Hyb}(Y, \text{STV}_k)$ is NP-hard, where Y is one of the following rules: plurality, Borda, maximin or Cup. Surprisingly, there are hybrid rules where the control by tie-breaking problem has a different complexity to the manipulation problem. In particular, $\textsc{Hyb}(plurality_k, plurality)$ is vulnerable to manipulation [Elkind and Lipmaa, 2005]. However, this hybrid is resistant to control by tie-breaking for unbounded k.

Theorem 17. *The control by tie-breaking problem for $\textsc{Hyb}(Plurality_k, Plurality)$ if k is unbounded is NP-complete.*

Proof: Our construction is similar to the construction in the proof of Theorem 3 in Elkind and Lipmaa [2005]. We reduce from an instance of the X3C problem where each item occurs in at most 3 subsets. We are given a set of items $A = \{a_1, \ldots, a_n\}$ with $|A| = n$ and subsets $S_1, S_2, \ldots, S_m \subset A$ with $|S_i| = 3$ for $i = 1, \ldots, m$. The question is whether there exists an index set I with $|I| = n/3$ and $\bigcup_{i \in I} S_i = S$. We build an election with $n + m + 2$ candidates: $C = A \uplus S \uplus \{p, d\}$. We have n candidates $A = \{a_1, \ldots, a_n\}$ that encode items, m candidates $S = \{s_1, \ldots, s_m\}$ that encode sets, a dummy candidate d and the preferred candidate p. Let T be a constant greater than $3nm$.

We introduce the following two sets of votes, $P = P_1 \uplus P_2$. The first set P_1 contains the following votes: T votes ($p \succ C \setminus \{p\}$), for each a_i, $i = 1, \ldots, n$, we introduce $T - 2$ votes ($a_i \succ C \setminus \{a_i\}$), for each a_i, $i = 1, \ldots, n$, we introduce 3 votes ($S^i \succ C \setminus S^i$), where S^i is the set of sets such that $a_i \in S_j$, $s_j \in S^i$. We introduce 4 votes ($d \succ C \setminus \{d\}$). Next we build P_2. Let n_j be the number of first places occupied by s_j in P_1. We know that $n_j \leq 3$. We introduce $3 - n_j$ votes ($s_j \succ d \succ C \setminus \{s_j, d\}$). The rest of the votes are irrelevant.

The initial plurality scores of the candidates are: $score(p) = T$, $score(a_i) = T-2$, $i = 1, \ldots, n$, $score(s_j) = 3$, $j = 1, \ldots, m$ and $score(d) = 4$. We set $k = m - n/3$.

During the first $m - n/3$ rounds, $m - n/3$ candidates from S are eliminated and the tie-breaking rule decides which $m - n/3$ out of m candidates to eliminate as all m candidates in S are tied . If the remaining $n/3$ candidates in S do not correspond to a cover, then an uncovered item a_i gets 3 points resulting in a plurality score of $T + 1$. Hence, p loses. Therefore, using the tie-breaking rule we must ensure that the remaining candidates from S form a cover. Finally, if a valid cover is selected, the maximal plurality score of d after k rounds is $4 + 3m$, the maximal plurality score of any surviving s_j is 9, the maximal plurality score of a_i is $T - 2$ and the score of p is T. Hence, p wins iff there exists a cover. □

Finally, we observe that if k or $m - k$ is bounded then $\textsc{Hyb}(Plurality_k, Plurality)$ is polynomial.

Theorem 18. *The control by tie-breaking problem for $\textsc{Hyb}(Plurality_k, Plurality)$ is polynomial if k or $m - k$ is bounded.*

Proof: If k is bounded, we can try all $O(m^k)$ possible tie-breaking decisions about candidates to eliminate. Similarly, if $m - k$ is bounded, we can try all $O(m^{m-k})$ possible tie-breaking decisions about candidates to survive. □

P	NP-complete
scoring rules, Cup, Nanson, Copeland, maximin Bucklin, fallback	STV, Baldwin ranked pairs, Coombs

Table 1: Complexity of control by tie-breaking. The result for ranked pairs is due to [Brill and Fischer, 2012].

9 Conclusion

We have studied the computational complexity of the control by tie-breaking problem. When the chair is only asked to break ties once to choose between one of the co-winners, the problem is trivially polynomial. However, in multi-round elections like STV, where the chair may have to break multiple ties, we proved that this control problem can be NP-complete (see Table 1 for a summary of results). Our results contain several surprises. For example, whilst it is NP-hard to compute a manipulating vote for a multi-round rule like Nanson's rule, control by tie-breaking is polynomial. Similarly, with a two-stage voting rule, even though the chair might only be asked to break ties at most twice, it can be NP-hard to control such a rule by breaking ties. Of course, many of our results are worst-case and may not reflect the difficulty of manipulation in practice. A number of recent theoretical and empirical results suggest that manipulation can often be computationally easy on average (e.g. [Conitzer and Sandholm, 2006; Procaccia and Rosenschein, 2007; Xia and Conitzer, 2008a; Xia and Conitzer, 2008b; Friedgut *et al.*, 2008; Walsh, 2009; Walsh, 2010]). Hence this work should be seen as just the first step in understanding computational issues surrounding the control of elections by breaking ties.

Acknowledgments

The authors are supported by the Australian Governments Department of Broadband, Communications and the Digital Economy, the Australian Research Council and the Asian Office of Aerospace Research and Development through grant AOARD-124056.

References

[Aziz *et al.*, 2012] H. Aziz, S. Gaspers, N. Mattei, N. Narodytska, and T. Walsh Ties matter: complexity of manipulation when tie-breaking with a random vote. In *Proceedings of the Twenty-Seventh AAAI Conference on Artificial Intelligence*, To Appear, 2012.

[Bartholdi and Orlin, 1991] J.J. Bartholdi and J.B. Orlin. Single transferable vote resists strategic voting. *Social Choice and Welfare*, 8(4):341–354, 1991.

[Bartholdi *et al.*, 1989] J.J. Bartholdi, C.A. Tovey, and M.A. Trick. The computational difficulty of manipulating an election. *Social Choice and Welfare*, 6(3):227–241, 1989.

[Bartholdi *et al.*, 1992] J.J. Bartholdi, C.A. Tovey, and M.A. Trick. How hard is it to control an election. *Mathematical and Computer Modelling*, 16(8-9):27–40, 1992.

[Brill and Fischer, 2012] M. Brill and F.A. Fischer. The price of neutrality for the ranked pairs method. In *Proceedings of the Twenty-Sixth AAAI Conference on Artificial Intelligence*, pages 1299–1305, 2012.

[Conitzer and Sandholm, 2003] V. Conitzer and T. Sandholm. Universal voting protocol tweaks to make manipulation hard. In *Proceedings of 18th IJCAI International Joint Conference on Artificial Intelligence*, pages 781–788, 2003.

[Conitzer and Sandholm, 2006] V. Conitzer and T. Sandholm. Nonexistence of voting rules that are usually hard to manipulate. In *Proceedings of the 21st National Conference on Artificial Intelligence*, pages 627–634, 2006.

[Conitzer *et al.*, 2007] V. Conitzer, T. Sandholm, and J. Lang. When are elections with few candidates hard to manipulate? *Journal of the ACM (JACM)*, 54(3):14, 2007.

[Conitzer *et al.*, 2009] V. Conitzer, M. Rognlie, and L. Xia. Preference functions that score rankings and maximum likelihood estimation. In *Proceedings of the Twenty-First International Joint Conference on Artificial Intelligence (IJCAI)*, pages 109–115, 2009.

[Davies *et al.*, 2012] J. Davies, N. Narodytska, and T. Walsh. Eliminating the weakest link: Making manipulation intractable? In *Proceedings of the Twenty-Sixth AAAI Conference on Artificial Intelligence* pages 1334-1339, 2012.

[Elkind and Lipmaa, 2005] E. Elkind and H. Lipmaa. Hybrid voting protocols and hardness of manipulation. In *Proceedings of 16th International Symposium on Algorithms and Computation (ISAAC 2005)*, volume 3827 of *Lecture Notes in Computer Science*, pages 206–215, 2005.

[Faliszewski *et al.*, 2008] P. Faliszewski, E. Hemaspaandra, and H. Schnoor. Copeland voting: ties matter. In *7th International Joint Conference on Autonomous Agents and Multiagent Systems*, pages 983–990, 2008.

[Faliszewski *et al.*, 2009] P. Faliszewski, E. Hemaspaandra, L.A. Hemaspaandra, and J. Rothe. Llull and copeland voting computationally resist bribery and constructive control. *Journal of Artificial Intelligence Research*, 35:275–341, 2009.

[Friedgut *et al.*, 2008] E. Friedgut, G. Kalai, and N. Nisan. Elections can be manipulated often. In *Proceedings of the 49th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2008)*, pages 243–249, 2008.

[McGarvey, 1953] D. McGarvey 1953. A theorem on the construction of voting paradoxes. *Econometrica* 21(4):608–610.

[Narodytska *et al.*, 2011] N. Narodytska, T. Walsh, and L. Xia. Manipulation of Nanson's and Baldwin's rules. In *Proceedings of the Twenty-Fifth AAAI Conference on Artificial Intelligence*, pages 713-718, 2011.

[Obraztsova and Elkind, 2011] S. Obraztsova and E. Elkind. On the complexity of voting manipulation under randomized tie-breaking. In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI 2011)*, pages 319–324, 2011.

[Obraztsova *et al.*, 2011] S. Obraztsova, E. Elkind, and N. Hazon. Ties matter: complexity of voting manipulation revisited. In *10th International Conference on Autonomous Agents and Multiagent Systems*, pages 71–78, 2011.

[Procaccia and Rosenschein, 2007] A. D. Procaccia and J. S. Rosenschein. Junta distributions and the average-case complexity of manipulating elections. *Journal of Artificial Intelligence Research*, 28:157–181, 2007.

[Rothe and Schend, 2012] J. Rothe and L. Schend. Control Complexity in Bucklin, Fallback, and Plurality Voting: An Experimental Approach. In *Proceedings of 11th International Symposium on Experimental Algorithms (SEA 2012)*, volume 7276 of *Lecture Notes in Computer Science*, pages 356–368, 2012.

[Tideman, 1987] N. Tideman. Independence of clones as a criterion for voting rules. *Social Choice and Welfare*, 4(3):185–206, 1987.

[Walsh, 2009] T. Walsh. Where are the really hard manipulation problems? The phase transition in manipulating the veto rule. In *Proceedings of the 21st International Joint Conference on Artificial Intelligence*, pages 324–329, 2009.

[Walsh, 2010] T. Walsh. An empirical study of the manipulability of single transferable voting. In *Proc. of the 19th European Conference on Artificial Intelligence (ECAI-2010)*, volume 215 of *Frontiers in Artificial Intelligence and Applications*, pages 257–262, 2010.

[Xia and Conitzer, 2008a] L. Xia and V. Conitzer. Generalized scoring rules and the frequency of coalitional manipulability. In *EC '08: Proceedings of the 9th ACM conference on Electronic Commerce*, pages 109–118, 2008.

[Xia and Conitzer, 2008b] L. Xia and V. Conitzer. A sufficient condition for voting rules to be frequently manipulable. In Lance Fortnow, John Riedl, and Tuomas Sandholm, editors, *Proceedings of the 9th ACM conference on Electronic Commerce (EC' 08)*, pages 99–108, 2008.